

IT 认证电子书



质 量 更 高 服 务 更 好

半年免费升级服务

<http://www.itrenzheng.com>

Exam : **NS0-141**

Title : Network Appliance High
Availability Exam 7G

Version : Demo

1. Which of the following hardware components is NOT required for Fabric Attached MetroCluster?

- A. VI-MC cluster adapter
- B. FC-HBA
- C. LC/LC node-to-node switch cables
- D. Network interface card

Answer: C

2. A mirrored volume has a failed disk, but there are no available spares in the pool from which the degraded plex was built. Data ONTAP will:

- A. Choose spare disks to reconstruct from the opposite pool
- B. Warn you that there are no available spares from the proper pool
- C. "Halt"after 24 hours if spare disks are not added to the proper pool
- D. Stop all applications that are running

Answer: B

3. A mirrored volume, volX, is composed of twelve 36-Gigabyte disks: six in volX/plex0, six in volX/plex1. A drive fails in plex1, but pool 1 contains only 72-Gigabyte spare disks. Data ONTAP will do which of the following?

- A. Choose a 72-Gigabyte disk and downsize it
- B. Halt after 24 hours of running in degraded mode
- C. Alert you that there are no 36-Gigabyte disks and wait for one to be inserted
- D. Stop any applications that are running

Answer: A

4. Select the parameter for which the Cluster Configuration Checker does NOT test.

- A. Mismatched option settings
- B. Network interface configuration errors
- C. Different versions of Data ONTAP on the local and partner nodes
- D. Open firmware versions

Answer: D

5. Please click the exhibit button.

```
filer1> vol status -r volX
Volume volX (online, resyncing) (block checksums)
  Plex /volX/plex0 (online, normal, active)
    RAID group /volX/plex0/rg0 (normal)

  RAID Disk Device  HA      SHELF BAY CHAN  Used (MB/blks)  Phys (MB/blks)
  -----
  parity            2a.16  2a      1      0  FC:A  34000/69632000  34768/71205200
  data              2a.18  2a      1      2  FC:A  34000/69632000  34768/71205200

  Plex /volX/plex1 (online, normal, resyncing 29% completed)
    RAID group /volX/plex1/rg0 (level-0 resync in progress)

  RAID Disk Device  HA      SHELF BAY CHAN  Used (MB/blks)  Phys (MB/blks)
  -----
  parity            8a.17  8a      1      1  FC:A  34000/69632000  34768/71205200
  data              8a.16  8a      1      0  FC:A  34000/69632000  34190/70022840
```

In this example, volX contains three 36-Gigabyte disk drives. To create a mirror for volX, you issue the command to create the mirror: vol mirror volX.

You immediately realize that the disks used to create the mirror are 72 Gigabytes in size. You want to take the plex offline and destroy it. You enter the vol status -r command and it displays the output shown in the exhibit.

To successfully take volX/plex1 offline and destroy it, which of the following must be done?

- A. Reboot the storage appliance.
- B. Wait for the "resyncing" operation to complete.
- C. Take volX/plex1 offline and split it off.
- D. Rename volX/plex1.

Answer: B

6. You are troubleshooting a clustered failover problem at a customer site. The customer explains that both storage appliances were NFS mounted onto the Solaris host. You attempt to create a file on the NFS mount point for the inaccessible appliance, but you get the following error:

NFS server filer2 not responding still trying.

Symptoms:

- 1) When not in takeover mode, both storage appliances operate normally.
- 2) In takeover mode, the taken over node is not accessible on the network.

You perform a "cf giveback" and you can now create a file over NFS on both appliances. What should you do first to resolve this issue?

- A. Replace the cluster interconnect card and run diagnostics "all" to ensure proper system operation.
- B. Replace the NIC adapter.
- C. Use the cluster-config-checker.cgi to identify the problem and corrective action.
- D. Submit a new bug report.

Answer: C

7. Which of the following commands would you issue to perform a "clean" shutdown of a clustered system without allowing takeover to occur?

- A. halt
- B. halt -f
- C. halt -d
- D. sys-shutdown

Answer: B

8. A storage appliance in a clustered configuration can't access its partner's mailbox disks. When you enter the cf status command, you see the following:

```
filer2> cf status
```

```
filer1 is up, takeover disabled because of reason (partner mailbox disks not accessible or invalid)
```

What should you check first to resolve this issue?

- A. Check volumes of all disks
- B. Check NIC adapter
- C. Check that initiator is set to "target" mode
- D. Check for disk viability or FC-AL problems

Answer: D

9. If an error message indicates that the cluster interconnect adapter itself is faulty, what is your corrective

action?

- A. Check the cabling of the Channel A disk shelf loop and reseal and tighten any loose cables.
- B. Check the cabling of the Channel B disk shelf loop and reseal and tighten any loose cables.
- C. Re-cable the cluster interconnect adapter or replace the adapter if needed.
- D. Check for network activity.

Answer: C

10. To perform a motherboard replacement on a cluster node without initiating a takeover, Network Appliance recommends using the halt command to perform the replacement.

- A. True
- B. False

Answer: B

11. You cannot disable the SyncMirror license if one or more mirrored volumes exist.

- A. True
- B. False

Answer: A

12. You have determined that the failure to a MetroCluster is a disaster. Which of the following steps should you take?

- A. Ensure the surviving node is isolated from its partner.
- B. Force the surviving node to take over the functions of its partner.
- C. Remount the failed partner's volume.
- D. All of the above

Answer: D

13. Which of the following MetroCluster configuration procedures is recommended for re-establishing a mirrored volume that was in a level-0 resynchronization state when it failed?

- A. Rejoin the two volumes.
- B. Re-create the synchronous mirror.

- C. Turn on power to the cluster node at the disaster site.
- D. Determine which volumes are at the remote and disaster sites.

Answer: B

14. Which one of the following statements is true of the `cf forcetakeover -d` command?

- A. The surviving cluster node takes over the functions of the failed partner.
- B. Mirrored volumes are implicitly split.
- C. Both of the above
- D. None of the above

Answer: C

15. When using MetroCluster, identify TWO methods for restricting access to the disaster site node.

- A. Use manual fencing.
- B. Turn off power to the disaster site node.
- C. Use the `cf forcetakeover -d` command.
- D. Use the `cf giveback -f` command.

Answer: AB

16. What is the maximum distance between a standard clustered pair?

- A. 10 meters
- B. 50 meters
- C. 100 meters
- D. 500 meters

Answer: D

17. Please click the "Exhibit" button.

19. Which TWO commands could be used if the cf giveback command fails?

- A. cf giveback -force
- B. cf giveback -f
- C. cf forcegiveback -f
- D. cf forcegiveback

Answer: BD

20. In cluster takeover mode, the shared interface assumes the identity of the _____.

- A. local interface only
- B. failed partner interface only
- C. local and failed partner interfaces
- D. None of the above

Answer: C

21. Which of the following is NOT correct when creating a SyncMirror volume?

- A. The number of disks added must be an even number and equally divided between the two plexes.
- B. Disks selected for Plex 0 and Plex 1 must come from the same disk pools.
- C. The number of disks selected for one plex must equal the number of disks selected from the other plex.
- D. Disks are selected first on the basis of equivalent bytes per sector (bps) size, then on the basis of the size of the disk. If there is no equivalent-sized disk, Data ONTAP takes a larger capacity disk and downsizes it.

Answer: B

22. What does the -n option in vol mirror volname -n specify when creating a SyncMirror volume?

- A. Shows the number of disks and creates the mirrored volume
- B. Specifies the existing volume using disks you select
- C. Shows the number of disks without creating the mirrored volume
- D. Consists of disk IDs of one or more available disks

Answer: C

23. Normally, before you split a SyncMirror volume, you ensure that both plexes are _____ and _____.

- A. plexes, pools
- B. online, operational
- C. offline, operational
- D. snapshots, contain the most recent snapshot

Answer: B

24. After issuing the vol split vol0/plex0 volNew command:

- A. There are two unmirrored volumes, vol0 and volNew.
- B. There are two new plexes, plex0 and plexNew.
- C. There is a new mirrored volume called volNew.
- D. There is a new plex called plexNew.

Answer: A

25. What communicates the state of a clustered partner?

- A. Status update
- B. Heartbeat
- C. Takeover command
- D. Cluster report

Answer: B

26. If the local node has its FC HBA connected to switch Bank 2, it owns the disks connected to switch bank ____.

- A. 1
- B. 3
- C. 4
- D. 2

Answer: A

27. Both appliances in a clustered configuration must have _____.

- A. one cluster license only
- B. a special set of licenses to allow takeover
- C. SyncMirror licenses
- D. the same license(s)

Answer: D

28. When cabling a standard storage appliance for SyncMirror, disk pool 0 volumes must be cabled to Fibre Channel adapter(s) located in what range of slots on the FAS960 appliance?

- A. Slots 3 through 8
- B. Slots 2 through 6
- C. Slots 3 through 11
- D. Slots 7 through 8

Answer: B

29. SyncMirror requires _____ license(s).

- A. syncmirror_local
- B. syncmirror_local and cluster_local
- C. syncmirror_remote
- D. syncmirror_remote and cluster_remote

Answer: A

30. A mirrored volume, volX, is composed of twelve 36-Gigabyte disks: six in volX/plex0, six in volX/plex1. A drive fails in plex1, but pool 1 contains only 72-Gigabyte spare disks. Data ONTAP will do which of the following?

- A. Choose a 72-Gigabyte disk and downsize it
- B. Halt after 24 hours of running in degraded mode
- C. Alert you that there are no 36-Gigabyte disks and wait for one to be inserted
- D. Stop any applications that are running

Answer: A